

CLIMBING MOUNT IMPROBABLE

Chapter 1: Facing Mount Rushmore

I have just listened to a lecture in which the topic for discussion was the fig. Not a botanical lecture, a literary one. We got the fig in literature, the fig as metaphor, changing perceptions of the fig, the fig as emblem of pudenda and the fig leaf as modest concealer of them, 'fig' as an insult, the social construction of the fig, D.H. Lawrence on how to eat a fig in society, 'reading fig' and, I rather think, 'the fig as text.'. The speaker's final *pensée* was the following. He recalled to us the Genesis story of Eve's tempting Adam to eat of the fruit of the tree of knowledge. Genesis doesn't specify, he reminded us, which fruit it was. Traditionally, people take it to be an apple. The lecturer suspected that actually it was a fig, and with this piquant little shaft he ended his talk.

This kind of thing is the stock-in-trade of a certain kind of literary mind, but it provokes *me* to literal-mindedness. The speaker obviously knows that there never was a Garden of Eden, never a tree of knowledge of good and evil. So what is he actually trying to say? I suppose he had a vague feeling that 'somehow', 'if you will', 'at some level', 'in some sense', 'if I may put it this way' it is somehow 'right' that the fruit in the story 'should' have been a fig. But enough of this foolery. It is not that we should be literalist and Gradgrindian, but our elegant lecturer was *missing* so much. There is genuine paradox and real poetry lurking in the fig, with subtleties to exercise an inquiring mind and wonders to uplift an aesthetic one. In this book I want to move to a position where I can tell the true story of the fig. But the fig story is only one out of millions that all have the same Darwinian grammar and logic – albeit the fig story is among the most satisfyingly intricate in all evolution. To anticipate the central metaphor of the book, the fig tree stands atop one of the highest peaks on the massif of Mount Improbable. But peaks as high the fig's are best conquered at the end of the expedition. Before that there is

much that needs to be said, a whole vision of life that needs to be developed and explained, puzzles that need to be solved and paradoxes that must be disarmed.

As I said, the story of the fig is, at the deepest level, the same story as for every other living creature on this planet. Though they differ in surface detail, all are variations on the theme of DNA and the thirty million ways by which it propagates itself. On our route we shall have occasion to look at spider webs – at the bewildering, though unconscious, ingenuity with which they are made and how they work. We shall reconstruct the slow, gradual evolution of wings and of elephant trunks. We shall see that ‘the’ eye, legendarily difficult though its evolution sometimes seems, has actually evolved at least 40 and probably 60 times independently all around the animal kingdom. We shall program computers to assist our imagination in moving easily through a gigantic museum of all the countless creatures that have ever lived and died, and their even more numerous, imaginary cousins who have never been born. We shall wander the paths of Mount Improbable, admiring its vertical precipices from afar, but always restlessly seeking the gently graded slopes on the other side. The meaning of the parable of Mount Improbable will be made clear, and much else besides. I need to begin by clarifying the problem of apparent design in nature, its relationship to true, human design and its relationship to chance. This is the purpose of Chapter 1.

The Natural History Museum in London has a quirky collection of stones that chance to resemble familiar objects: a boot, a hand, a baby’s skull, a duck, a fish. They were sent in by people who genuinely suspected that the resemblance might mean something. But ordinary stones weather into such a welter of shapes, it is not surprising if occasionally we find one that calls to mind a boot, or a duck. Out of all the stones that people notice as they walk about, the museum has preserved the ones that they pick up and keep as

curiosities. Thousands of stones remain uncollected because they are just stones. The coincidences of resemblance in this museum collection are meaningless, though amusing. The same is true when we think we see faces, or animal shapes, in clouds or cliff profiles. The resemblances are accidents.

Fig 1.1 A pure accident. President Kennedy's face in a hillside.

This craggy hillside is supposed to suggest the profile of the late President Kennedy. Once you have been told, you can just see a slight resemblance to either John or Robert Kennedy. But some don't see it at all, and it is certainly easy to believe that the resemblance is accidental. You couldn't, on the other hand, persuade a reasonable person that Mount Rushmore, in South Dakota, had just happened to weather into the features of Presidents Washington, Jefferson, Lincoln and Theodore Roosevelt. We do not need to be told that these were deliberately carved (under the direction of Gutzon Borglum). They are obviously not accidental: they have design written all over them.

The difference between Mount Rushmore and the weathered likeness of John Kennedy (or Mont St Pierre in Mauritius or all such curiosities of natural weathering) is this. The sheer number of details in which the Mount Rushmore faces resemble the real thing is too great to have come about by chance. The faces are clearly recognizable, moreover, when seen from all different angles. Figure 1.1's chance resemblance to President Kennedy, on the other hand, is only noticed if the cliff is seen from a particular angle and in a particular light. Yes, a rock can weather into the shape of a nose seen from a certain vantage point, and maybe a couple of other rocks happen to have tumbled into the shape of lips. It is not much to ask of chance that it should produce a modest coincidence like this, especially if viewers have all possible angles to choose from and only one gives the resemblance (and there is the added fact, which I'll return to in a

moment, that the human brain seems actively eager to see faces: it seeks them out). But Mount Rushmore is another matter. Its four heads are clearly *designed*. A sculptor conceived them, drew them out on paper, made meticulous measurements all over the cliff, and supervised teams of workmen who wielded pneumatic drills and dynamite to carve out the four faces, each 60 feet high. The weather *could* have done the same job as the artfully deployed dynamite. But of all the possible ways of weathering a mountain, only a tiny minority would be speaking likenesses of four particular human beings. Even if we didn't know the history of Mount Rushmore, we'd estimate the odds against its four heads being carved by accidental weathering as astronomically high – like tossing a coin 40 times and getting heads every time.

I think that the distinction between accident and design is clear, in principle if not always in practice, but this chapter will introduce a third category of objects which is harder to distinguish. I shall call them *designoid*, pronounced 'design-oid' not 'dezzignoid.' Designoid objects are living bodies and their products. Designoid objects *look* designed, so much so that some people – probably, alas, most people – think that they *are* designed. These people are wrong. But they are right in their conviction that designoid objects cannot be the result of chance. Designoid objects are not accidental. They have in fact been shaped by a magnificently nonrandom process which creates an almost perfect illusion of design.

Figure 1.2 An undesigned, but not accidental, resemblance. An ant (a) and an ant-mimicking beetle (b).

Here is a living sculpture. Beetles in general don't look like ants. So, if I see a beetle that looks almost exactly like an ant – a beetle, moreover, that makes its living entirely in an ant's nest – I shall rightly suspect that the coincidence means something. The top animal is actually a beetle – its closer cousins are common or garden beetles – but it looks like an ant, walks like an ant, and lives

among ants in an ants' nest. The one at the bottom is a real ant. As with any realistic statue, the resemblance to the model is not an accident. It demands an explanation other than sheer chance. What kind of an explanation? Since all beetles that look strikingly like ants live in ants' nests, or at least in close association with ants, could it be some chemical substance from the ants, or some infection from the ants, rubbing off on the beetles and changing the way they grow? Definitely not. We shall come to the true explanation – Darwinian natural selection – later. For the moment, it is enough that we are sure this resemblance, and other examples of 'mimicry,' are not accidental. They are either designed or they are due to some process that produces results just as impressive as design. We shall look at some other examples of animal mimicry, leaving open, for the moment, the explanation of how these remarkable resemblances come about.

Figure 1.3 a. A real termite. b. a beetle mimicking a termite. c. How the trick is done

The previous example shows what a good job beetle flesh can do if it 'sets out to mimic' a different kind of insect. But now look at the creature in Figure 1.3b. It appears to be a termite. Figure 1.3a is a real termite, for comparison. The specimen in Figure 1.3b is an insect, but it is not a termite. It is, in fact, a beetle. I admit that I've seen better mimics in the insect world, including the ant-mimicking beetle of the previous example. The 'beetle' here is just a little odd. Its legs seem to lack proper joints, like little twisty balloons. Since, like any other insect, a beetle has jointed legs at its disposal, you might hope for a better shot at mimicking a termite's jointed legs. So, what is the solution to this conundrum? Why does this statue look like an inflated dummy rather than like a real, jointed insect. The answer can be seen in Figure 1.3c, which is one of the most astonishing spectacles in all natural history. It shows the termite-mimicking beetle in side view. The true head of the beetle is a diminutive affair (you can see the eye just near the normal, jointed

antennae), attached to a slender trunk or thorax bearing three normal, jointed beetle legs, on which it actually walks. It is with the abdomen that the trick is done. It is arched backwards so that it hangs over and completely covers the head, thorax and legs like a parasol. The entire 'termite' is constructed from the (anatomically) rear half of the beetle's abdomen. The 'termite head' is the rear tip of the beetle's abdomen. The 'termite legs' and 'antennae' are flapping excrescences of the abdomen. No wonder the quality of the mimicry is not quite up to the standard of the beetle's ant-mimicking cousin of the previous picture. This termite-mimicking beetle, by the way, lives in termite nests, making its living as a parasite in much the same way as Figure 1.2's ant-mimicking beetle makes its living among ants. Although the quality of the resemblance is less, when you consider its starting materials the termite-mimicking beetle seems to achieve a more impressive feat of sculpture than the ant-mimicking beetle. This is because the ant-mimic does it by modifying each bit of its body to look like the corresponding bit of the ant's body. But the termite-mimic does it by modifying a completely different bit of itself - the abdomen - to look like all the bits of the termite.

Figure 1.4 Perfection of camouflage. Leafy Sea Dragon.

My own favourite among animal 'statues' is the leafy sea dragon. It is a fish, a kind of sea-horse, whose body is sculpted into the shape of seaweed. This gives it protection, for it lives among seaweed and is remarkably difficult to see there. Its mimicry is too uncannily good to be accidental in any simple sense. It lies closer to Mount Rushmore than to the Kennedy cliff. My confidence is based partly upon the sheer number of ways in which it impresses us by looking like something that it isn't; and partly upon the fact that fish don't normally have projections of anything like that shape. In this respect the leafy sea dragon's feat compares with the termite mimic, rather than the ant mimic.

So far we have talked of objects that impress us as realistic sculptures do, objects that we feel can't be accidental because they look too strikingly like other objects. Leafy sea dragons and ant-mimicking beetles are designoid statues: they overwhelmingly look as if they have been designed by an artist to resemble something else. But statues are only one kind of object that humans design. Other human artefacts impress us not by resembling something but by being unmistakably useful for some purpose. An aeroplane is useful for flying. A pot is useful for holding water. A knife is useful for cutting things.

If you offered a reward for stones that were naturally sharp enough to cut things, and also for stones that happened to be of a shape to hold water, you'd probably be sent some effective makeshifts. Flints often fracture in such a way as to leave a good keen edge, and if you wandered the quarries and screes of the world you'd certainly find some handy natural blades. Among the richness of shapes into which stones can weather, some would happen to include concavities that hold water. Certain types of crystal naturally encrust around a hollow, albeit chunky, sphere which, when it splits in half, yields two serviceable cups. These stones even have a name: geode. I use a geode as a paperweight on my desk, and I'd use it to drink from if its interior were not roughly pitted and therefore hard to wash. It is easy to devise measures of efficiency that would show up natural pots as less efficient than manmade ones. Efficiency is some measure of benefit divided by cost. The benefit of a pot could be measured as the quantity of water that it holds. Cost can conveniently be measured in equivalent units: the quantity of the material of the pot itself. Efficiency might be defined as the volume of water that a pot can hold divided by the volume of material that goes to make the pot itself. The hollow stone on my desk holds 87.5 cc of water. The volume of the stone itself (which I measured by Archimedes's famous Eureka-in-the-Bath method) is 130 cc. The efficiency of this 'pot' is therefore about two thirds. This is a very low

efficiency, not surprisingly so since the stone was never designed to hold water. It just happens to hold water. I have just done the same measurements on a wineglass, whose efficiency turns out to be about 3.5. My silver cream jug is even more efficient. It holds 250 cc of water while the silver of which it is made displaces a mere 20 cc. Its efficiency is therefore as high as 12.5.

Not all human-designed pots are efficient in this sense. A chunky pot from the kitchen cupboard holds 190 cc of water while using up a massive 400 cc of marble. Its 'efficiency' is therefore only 0.475, even lower than the totally undesigned hollow stone. How can this be? The answer is revealing. This marble pot is in fact a mortar. It is not designed just to hold liquid. It is a hand mill for grinding spices and other foods with a pestle: a stout rod which is wielded with great force against the inside of the mortar. You couldn't use a wineglass as a mortar: it would shatter under the force. The measure of efficiency that we devised for pots is not suitable when the pot is designed as a mortar. We should devise some other benefit/cost ratio, where benefit takes account of strength against being broken by a pestle. Would the natural geode, then, qualify as a well-designed mortar? It would probably pass the strength test but if you tried to use it as a mortar its rough and craggy interior would soon prove a disadvantage, the crevices protecting grains from the pestle. You'd have to improve your measure of the efficiency of a mortar by including some index of smoothness of internal curvature. That my marble mortar is designed can be discerned from other evidence: its perfectly circular plan section, coupled with its elegantly turned lip and plinth seen in elevation.

We could devise similar measures of the efficiency of knives, and I have no doubt that the naturally flaked flints that we happen to pick up in a quarry would compare unfavourably, not only with Sheffield steel blades but with the elegantly sculpted flints that museums display in Late Stone Age collections.

There is another sense in which natural, accidental, pots and knives are inefficient compared with their designed equivalents. In the course of finding one usefully sharp flint tool, or one usefully watertight stone vessel, a huge number of useless stones had to be examined and discarded. When we measure the water held by a pot, and divide by the volume of stone or clay in the material of the pot, it might be fairer to add into the denominator the cost of the stone or clay discarded. In the case of a manmade pot thrown on a wheel, this additional cost would be negligible. In the case of a carved sculpture the cost of discarded chippings would be present but small. In the case of the accidental, *objet trouvé* pot or knife, the 'discard cost' would be colossal. Most stones don't hold water and are not sharp. An industry that was entirely based upon *objets trouvés*, upon *found* objects as tools and utensils, rather than artificially shaped tools and utensils, would have a huge dead weight of inefficiency in the spoil heaps of alternatives discarded as useless. Design is efficient compared with finding.

Figure 1.5 A designoid pot. Pitcher plant.

Let's turn our attention now to designoid objects – living things that look as though they have been designed but have actually been put together by a completely different process – beginning with designoid pots. The pitcher plant could be seen as just another kind of pot, but it has an elegant 'economy ratio,' comparable to the wine glass that I measured, if not the silver jug. It gives every appearance of being excellently well designed, not just to hold water but to drown insects and digest them. It concocts a subtle perfume which insects find irresistible. The smell, abetted by a seductive colour pattern, lures prey to the top of the pitcher. There the insects find themselves on a steep slide whose treacherous slipperiness is more than accidental, set about with downward-facing hairs well-placed to impede their last struggle. When they fall, as they nearly always do,

into the dark belly of the pitcher, they find more than just water in which to drown. The details, brought to my attention by my colleague Dr Barrie Juniper, are remarkable and I'll briefly tell the story.

It is one thing to trap insects but the pitcher plant lacks jaws, muscles and teeth with which to reduce them to a state fit for digesting. Perhaps plants could grow teeth and munching jaws but in practice there is an easier solution. The water in the pitcher is home to a rich community of maggots and other creatures. They live nowhere else but in the enclosed ponds created by pitcher plants, and they are endowed with the jaws that the plant itself lacks. The corpses of the pitcher plant's drowned victims are devoured and decomposed by the mouthparts and digestive juices of its maggot accomplices. The plant itself subsists on the detritus and excretory products, which it absorbs through the lining of the pitcher.

The pitcher plant doesn't just passively accept the services of maggots that happen to fall into its private pool. The plant works actively to provide the maggots with a service that they need in their turn. Analyse the water in a pitcher plant and you find a singular fact. It is not foetid as might be expected of standing water in such conditions, but strangely rich in oxygen. Without this oxygen the vital maggots could not flourish, but where does it come from? It is manufactured by the pitcher plant itself, and the plant gives every apparent indication of being specifically designed to oxygenate the water. The cells that line the pitcher are richer in oxygen-producing chlorophyll than the outside cells that face the sun and air. This surprising reversal of apparent common sense is explicable: the inside cells are specialised to secrete oxygen directly into the water inside the pitcher. The pitcher plant does not just borrow its vicarious jaws: it hires them, paying in the currency of oxygen.

Other designoid traps are common. The Venus's fly trap is as

elegant as the pitcher plant, with the added refinement of moving parts. The insect prey releases the trap by triggering sensitive hairs on the plant, whose jaws smartly close. The spider web is the most familiar of all animal traps, and we shall do it justice in the next chapter. An underwater equivalent is the net constructed by stream-dwelling caddis fly larvae. Caddis larvae are also notable for their feats as builders of houses for themselves. Different species use stones, sticks, leaves or tiny snail shells. A familiar sight in various parts of the world is the conical trap of the ant lion. This fearsome creature is the larva of – what could sound more gentle? – a lacewing fly. The ant-lion lurks just under the sand at the bottom of its pit, waiting for ants or other insects to fall in. The pit achieves its almost perfectly conical shape – which makes it hard for victims to claw their way out – not by design but as a consequence of some simple rules of physics, exploited by the way the ant-lion digs. From the bottom of the descending pit, it flicks sand right over the edge with a jerk of the head. Flicking sand from the bottom of a pit has the same effect as draining an hourglass from below: the sand forms itself naturally into a perfect cone of predictable steepness.

Figure 1.6 Designoid pots made by animal artisans. (a) Potter wasp and (b) mason bee pots

Figure 1.6 brings us back to pots. Many solitary wasps lay their eggs on prey, stung to paralysis and hidden in a hole. They seal the hole up so that it is invisible, the larva feeds on the prey inside and finally emerges as a winged adult to complete the cycle. Most species of solitary wasp dig their nesting hole in the ground. The potter wasp makes its own ‘hole’ out of clay – a round pot, up a tree, mounted inconspicuously on a twig. Like the pitcher plant, this pot would score favourably on our efficiency test for apparent design. Solitary bees show a similar pattern of nesting in holes, but they feed their larvae on pollen instead of animal prey. Like the potter among wasps, many species of mason bee build their own pot nest. The pot

in Figure 1.6 is made not from clay but from small stones cemented together. Apart from its resemblance to an efficient, man-made receptacle, there is something else rather wonderful about the particular specimen photographed. You see only one pot here, but there are actually four. The other three have been covered by the bee with hardened mud, to give an exquisite match to the surrounding rock. No predator would ever find the young larvae growing up in the pots. The only reason this cluster was seen, by my colleague Christopher O'Toole on a visit to Israel, is that the bee had not quite finished covering the last pot.

These insect pots have all the hallmarks of 'design.' In this case, unlike the pitcher plant, they really were fashioned by the actions of a skilled – albeit probably unconsciously so – creature. The pots of the potter wasp and mason bee seem, on the face of it, closer to human-made pots than to the pitcher plant. But the wasp and the bee didn't consciously or deliberately design their pots. Although they were shaped, out of clay or stones, by behavioural actions of the insects, this is not importantly different from the way the insects' own bodies were made during embryonic development. This may sound odd but let me explain. The nervous system grows in such a way that the muscles and limbs and jaws of the living wasp move in certain coordinated patterns. The consequence of these particular clockwork limb movements is that clay is gathered and fashioned into the shape of a pot. The insect very probably does not know what it is doing, nor why it is doing it. It has no concept of a pot as a work of art, or as a container, or as a brood chamber. Its muscles just move in the way its nerves dictate, and a pot is the result. So for this reason we firmly – if wonderingly – classify the wasp and bee pots as designoid not designed: not shaped by the animal's own creative volition. Actually, to be fair, I cannot know for certain that wasps lack creative volition and true design. It is enough for me that my explanation works even if they do. The same goes for birds' nests and bowers, caddis houses and caddis nets, but not for the

sculptures of Mount Rushmore or the tools used to hew them – they really are designed.

Figure 1.7 Designoid craftsmanship. (a) Weaver Bird Nest and (b) Tailor bird with its nest

Karl von Frisch, the famous Austrian zoologist who deciphered the bee dance, once wrote: “If we imagined for a moment that termites were as tall as human beings, their tallest hillocks, enlarged on the same scale, would be nearly a mile high, four times the height of New York’s Empire State Building.”

Figure 1.8 Insect skyscrapers. Compass Termite Mounds in Australia

The skyscrapers of Figure 1.8 were made by Australian compass termites. They are called compass termites because their mounds are always lined up North-South – they can be used as compasses by lost travellers (as can satellite dishes, by the way: in Britain they seem all to face south). The advantage of this for the termites is that the broad, flat surfaces of the mound are warmed by the early morning and late afternoon sun. But the mound is protected from the fierce noonday sun since only the sharp edge is presented to the north – where the sun is at midday in the Southern Hemisphere. The termites behave as if they had designed this clever trick themselves. But the principle by which their building behaviour appears intelligent is identical to the principle by which the jaws and legs of the termites appear designed. Neither of them is designed. Both are designoid.

Animal artefacts, like caddis and termite houses, birds’ nests or mason bee pots, are fascinating but they are a special case among designoid things – an intriguing curiosity. The name ‘designoid’ primarily refers to living bodies themselves and their parts. Living

bodies are put together not by skilled hands, beaks or jaws, but by the convoluted processes of embryonic growth. A mind addicted to exhaustive classification systems might see artefacts like wasp pots as 'second order designoid objects' or as an intermediate category between designed and designoid, but I think this is simply confusing. Admittedly the pot is made of mud not living cells, and it is shaped by limb movements which superficially resemble the hand movements of a human potter. But all the 'design', all the elegance, all the fitness of the pot to perform a useful task, come from very different sources in the two cases. The human pot is conceived and planned by a creative process of imagination in the head of the potter, or by deliberate imitation of the style of another potter. The wasp pot gets its elegance and fitness to its task from a very different process - from exactly the same process, indeed, as gave elegance and fitness to the wasp's own body. This will become clearer if we continue our discussion of living bodies as designoid objects.

One of the ways we recognize both true design and designoid pseudo-design is that we are impressed by resemblances between objects and other objects. The Mount Rushmore heads are obviously designed because they resemble real Presidents. The leafy sea dragon's resemblance to seaweed is equally obviously not an accident. But mimicry like this, and like the resemblance of beetle to termite, or stick insect to twig, is by no means the only kind of resemblance that impresses us in the living world. Often we are stunned by the resemblance between a living structure and a man-made device that does the same job. The 'mimicry' between human eye and manmade camera is too well known to need illustrating here. Engineers are often the people best qualified to analyse how animal and plant bodies work, because efficient mechanisms have to obey the same principles whether they are designed or designoid.

Figure 1.9 Animals with similar needs often resemble each

other more than they resemble their closer relatives. The Algerian hedgehog (a) *Erinaceus algirus* is a close cousin of the shrew hedgehog (b) *Neotetracus sinensis*. The greater hedgehog tenrec (c) *Setifer setosus* is a close cousin of the long-tailed tenrec (d) *Microgale melanorrhachis*.

Often living bodies have converged upon the same shape as each other, not because they are mimicking each other but because the shape that they share is separately useful to each of them. The hedgehog and the spiny tenrec in Figure 1.9 are so similar that it seemed almost a waste of effort to draw both of them. They are moderately closely related to each other, both being members of the Order Insectivora. Yet other evidence shows that they are sufficiently unrelated for us to be sure that they evolved their prickly appearance independently, presumably for parallel reasons: prickles give protection from predators. Each of the prickly animals is pictured next door to a shrew-like animal which is a closer cousin to it than the other prickly animal is. Figure 1.10 gives another example. Animals that swim fast near the surface of the sea often converge on the same shape as each other. It is the shape that engineers would recognize as streamlined. The picture shows a dolphin (mammal), an extinct ichthyosaur (which we can think of as the reptilian equivalent of a dolphin), a marlin (bony fish) and a penguin (bird). This kind of thing is called convergent evolution.

Figure 1.10 Convergent evolution: Independently evolved streamlining. Top to bottom: Dolphin, ichthyosaur, fish, penguin.

Apparent convergence is not always so meaningful. Those people – not all of them missionaries – that dignify face to face copulation as diagnostic of higher humanity may be charmed by the millipedes in Figure 1.11. If we call this convergence, it is probably not due to convergent needs: rather, there are only so many ways in which a

male and a female can juxtapose their bodies, and there could be lots of reasons for hitting upon any one of them.

Figure 1.11 Copulating millipedes, Coco-de-mer and Death's Head Hawk Moth.

This brings us full circle to our opening topic of pure accident. There are some living things that resemble other objects but where the resemblance is probably not strong enough to be anything other than accidental. The Bleeding Heart pigeon has a tuft of red feathers, so placed as to create the illusion of a mortal wound to the breast, but the resemblance is unlikely to mean anything. Equally accidental is the coco-de-mer's resemblance to a woman's loins (Figure 1.11). As in the case of the Kennedy profile in the cliff, the reason for feeling that these resemblances are coincidence is statistical: the sort of point we were making in the previous chapter about 'numbers of ways in which things could happen.' The pigeon's bleeding heart amounts only to a gash of red feathers. The coco-de-mer's apparent 'mimicry' is admittedly impressive. It involves two or three features, not just one. It even has a suggestion of pubic hair. But the human brain works hard and actively seeks resemblances, especially to bits of our bodies that we find particularly interesting. I suspect that this is going on in our perception of the coco-de-mer, just as it is in our recognition of the Kennedy hillside.

The same goes for the Death's Head Hawk moth. Indeed our brain has an almost indecent eagerness to see faces, which is the basis of one of the most striking illusions known to psychologists. If you get an ordinary face mask from a fancy dress shop and hold it up, hollow side facing another person (with a background that makes the eye holes conspicuous), that person is likely to see it standing out as a solid face. This has a very odd consequence which you'll discover if you gently rotate the mask from side to side. Remember that the viewer's brain 'thinks' it is a solid face, but the object is actually a

hollow mask. When the hollow mask moves to the left, the only way to reconcile the eyes' report with the brain's presumption that the face is solid, is to suppose that it moves in the opposite direction. And this is exactly the illusion that the viewer will see. The face will appear to be rotating in a direction opposite to the actual direction of rotation. There is much more that could be said about this fascinating illusion, but here I am just using it as evidence of the eagerness of the human brain to see faces, when given the smallest encouragement.

So it seems quite likely that the Death's Head Hawk moth resembles a face by accident, abetted by the human brain's innate eagerness to see faces. I should add, however, that one of our most respected evolutionary theorists, Robert Trivers, now of Rutgers University, New Jersey, believes that face mimicry on the backs of insects can be an adaptation to scare would-be predators such as birds (we think of the face on the moth as a human skull, but it could equally be a monkey's face). He could well be right, in which case I should have placed the example under my 'designoid' heading. For a different reason, the same may be true of another apparent face mimic, the Japanese Samurai crab. This crab has on its back a likeness (not, I have to say, a stunningly impressive likeness) of the fierce features of a samurai warrior. The suggestion has been made that over the centuries Japanese fishermen, abetted by the human brain's natural eagerness to see faces, have noticed a slight resemblance to a face on the backs of some individual crabs. For reasons of superstition or respect, fishermen did not wish to kill crabs with a human-like face (maybe especially a samurai-like face) on its back, so they threw them back into the sea instead. Many a crab's life was saved, according to this theory, by its humanoid face, and those crabs with the most pronounced human features in any one generation contributed a disproportionate share of offspring to the next generation. Later generations therefore had a head start over earlier ones, and the resemblance gradually increased. Whether or not

these apparent face mimics are accidental or designoid, I presume that the coco de mer's apparent mimicry is a pure accident, not enhanced by subsequent selection.

When we were discussing how to acquire a stone knife simply by finding it, we agreed that you could 'make' a sharp knife by examining all the stones in the world and discarding the blunt ones – the great majority. If you searched enough screes and quarries, you'd surely find a stone with not only a keen blade but a convenient handle as well. It is only a partial oversimplification to say that the pharmaceutical industry works by examining lots of molecules generated at hazard and then testing the efficacy of the minority that seem promising. But we agreed that *finding*, as a method of acquiring a useful implement, was supremely inefficient. Far better to take a suitable material like stone or steel and hone it or sculpt it by design. Yet this is not how designoid objects – living things bearing the illusion of design – are made. Living things come about ultimately by a process that is rather more like 'finding', but it is different from pure finding in a very significant respect. It may seem an odd fact to notice of a stone, but I'll mention it anyway and follow where it leads. A stone does not have children. If stones had children like themselves, those children would inherit from their parents the attribute of having children. This implies grandchildren and great grandchildren for unspecified generations. A far-fetched speculation it might be thought and, in any case, so what? To answer this, turn to something whose sharpness may be equally inadvertent but which does have offspring.

The hard, strap-like leaves of some reeds have quite sharp edges. This sharpness is probably an incidental byproduct of other properties of the leaf. You can cut yourself on a reed, enough to annoy but not deeply enough for the sharpness to provoke suspicion of design. No doubt some leaves are sharper than others and you could search the lake shore for the sharpest reed you could find.

Now here is where we part company with stones. Don't just cut with your reed knife, *breed from it*. Or breed from the same plant from which you plucked it. Allow the sharpest plants to cross-pollinate, kill the blunt plants: it doesn't matter how you do it, just see to it that the sharpest plants do most of the breeding. Not just once, but generation after generation. As the generations pass you'll notice that there are still bluntish reeds and sharpish reeds around, but the average reed will become steadily sharper. After 100 generations you'd probably have bred something that would give you a decent close shave. If you bred for rigidity at the same time as breeding for keenness of edge, you could eventually cut your throat with a broken reed.

In a sense you have done no more than *find* the quality you seek: no carving, whittling, moulding or whetstone grinding, just finding the best of what is already there. Sharp leaves have been found, blunt leaves have been discarded. It is like the story of finding sharp stones but with one significant addition: the process is cumulative. Stones don't breed whereas leaves, or rather the plants that make leaves, do. Having found the best blade of a generation you don't just use it until it wears out. You ratchet your gain by breeding from it, transmitting its virtue to the future where it can be built upon. This process is cumulative and never-ending. You are still just finding and finding, but because genetics enables cumulative gain the best specimen you can find in a late generation is better than the best you can find in an earlier generation. This, as we shall see in the next chapter, is what Climbing Mount Improbable means.

Figure 1.12 The power of artificial selection to shape plants. All these domestic vegetables have been bred by humans from the same ancestor, the wild cabbage *Brassica oleracea*.

The steadily sharpening reed was an invention to make a point. There are, of course, real examples of the same principle at work. All

the plants in Figure 1.12 are derived from one wild species, the wild cabbage, *Brassica oleracea*. It is a rather nondescript plant which doesn't look much like a cabbage. Humans have taken this wild plant and, over a short period of centuries, shaped it into these really very different kinds of food plants. It is a similar story with dogs.

Figure 1.13 The power of artificial selection to shape animals. All these domestic dogs have been bred by humans from the same wild ancestor, a wolf (top left): Chihuahua, long-haired dachshund, whippet, English bulldog, great Dane.

Although hybrids between dogs and jackals and between dogs and coyotes do occur, it is now accepted by most authorities that all breeds of domestic dogs are descended from a wolf ancestor (top left) who lived maybe a few thousand years ago. It is as though we humans had taken wolf flesh and shaped it like a clay pot. But we didn't, of course, literally knead and press wolf flesh into the shape of, say, a whippet or a dachshund. We did it by cumulative finding, or, as it is more conventionally put, selective breeding or artificial selection. Whippet-breeders *found* individuals that looked a little bit more whippet-like than the average. They bred from them, and then found the most whippet-like individuals of the next generation, and so on. Of course it wouldn't have been so relentlessly simple as that, and the breeders wouldn't have had the concept of a modern whippet in their heads as a distant target. Perhaps they just liked the look of the physical characteristics that we would now recognize as whippet-like, or perhaps those visible qualities came along as a byproduct of breeding for something else, like proficiency in hunting rabbits. But whippets and dachshunds, Great Danes and bulldogs, were made by a process that resembled finding more than it resembled clay-modelling. Yet it still is not the same as pure finding, because it is cumulative over generations. That is why I called it *cumulative finding*.

Accidental objects are simply found. Designed objects aren't found at all, they are shaped, moulded, kneaded, assembled, put together, carved: in one way or another the individual object is pushed into shape. Designoid objects are cumulatively found, either by humans as in the case of domestic dogs and cabbages, or by nature in the case of, say, sharks. The fact of heredity sees to it that the *accidental* improvements *found* in each generation are accumulated over many generations. At the end of many generations of cumulative finding, a designoid object is produced which may make us gasp with admiration at the perfection of its apparent design. But it is not real design, because it has been arrived at by a completely different process.

It would be nice to be able to demonstrate this process whenever we wish. Dog generation times are a bit shorter than ours but, even so, it takes more than a human lifetime to propel dog evolution to any notable extent. You can speed the process no end with a computer. Humans have bred Chihuahuas in about a ten thousandth of the time it took nature to breed wolves from their Chihuahua-sized (though not Chihuahua-shaped), insectivorous ancestors who lived back when the dinosaurs died out. Even so, artificial selection of real, living creatures – at least creatures bigger than bacteria – is too slow to make an impressive demonstration for impatient and short-lived humans. Computers, whatever their faults, are blindingly fast, and they can simulate anything that can be precisely defined, which includes reproductive processes like those of animals and plants. If you simulate heredity, that most basic condition for life, and provide for occasional random mutation, it is truly startling what can evolve before your eyes in a few hundred generations of selective breeding. I pioneered this approach in my book *The Blind Watchmaker* using a computer program of the same name. With this program you can breed, by artificial selection, creatures called computer biomorphs. Computer biomorphs are all bred from a common ancestor that looks

like this , in very much the same sense as all the breeds of dogs were bred from a wolf. Litters of progeny with random 'genetic mutations' appeared on the computer screen and a human chose which member of each litter to breed from. This needs some explanation. First, what does it mean to speak of 'progeny', of 'genes,' and of 'mutations' in the case of these computer objects? All biomorphs have the same kind of 'embryology.' All are basically built as a branching tree, or a segmental series of such trees joined up to one another. Details of the tree(s), such as how many branches there are, and the lengths and angles of the various branches, are controlled by 'genes' which are just numbers in the computer. Genes in real trees, like genes in us and genes in bacteria, are coded messages written in the language of DNA. The DNA is copied from generation to generation with great, though not perfect, fidelity. Within each generation, the DNA is 'read out' and has an influence on the shape of the animal or plant. Figure 1.14 shows how, in real trees and in computer biomorph trees, changes in just a few genes can alter the shape of the whole plant by changing the programmed rules of growth as each new twig is budded off. Biomorph genes are not made of DNA but this difference is trivial for our purposes. DNA is digitally coded information just like numbers in a computer, and numerical 'genes' pass down the generations of biomorphs in the same kind of way as DNA passes down generations of plants or animals.

Figure 1.14. Real trees and computer biomorph trees to show how varieties of the same species can vary in shape because of quite minor changes in growth rules. Several species of tree have a weeping variety and several species have converged upon the skypointing, 'Lombardy' variety.

When a biomorph has a child, the child inherits all the genes of its parent (it only has the one parent, for there is no sex) but with some possibility of random mutation. A mutation is a slight random increase or decrease in the numerical value of a gene. So a child

might be like its parent but with a slightly steeper angle to one of its branches because the numerical value of its Gene 6 has increased from 20 to 21. When in biomorph breeding mode, the computer draws a biomorph in the centre of the screen, surrounded by a litter of randomly mutated offspring. Because their genes have changed only slightly, the offspring always bear a family resemblance to the parent, and to each other, but they often display slight differences that the human eye can detect. Using the computer mouse, a human chooses one of the screenful of biomorphs for 'breeding.' The screen goes blank except for the chosen biomorph, which glides to the parental slot at the centre of the screen and then 'spawns' a new litter of mutant offspring around itself. As generations go by, the selector can guide evolution in very much the same way as humans guided the evolution of domestic dogs, but much faster. One of the things that surprised me when I first wrote the program was how quickly you could evolve away from the original tree shape. I found that I could home in on an 'insect' or a 'flower', a 'bat', a 'spider' or a 'spitfire.' Each one of the biomorphs in Figure 1.15 is the end product of hundreds of generations of breeding by artificial selection. Because the creatures breed in a computer, you can whistle through many generations of evolution in a matter of minutes. A few minutes of playing with this program on a modern, fast computer gives you a hands-on, vivid feeling for how Darwinian selection works. The biomorphs in the safari park of Figure 1.15 seem to me to resemble wasps, butterflies, spiders, scorpions, flatworms, lice, and other 'creatures' that look vaguely biological even if they don't resemble *particular* species on this planet. Yet all are cousins of the trees among which they stand, and of the squadron of 'spitfires' in the top right corner. They are pretty close cousins, at that. All have the same number of genes (16). They differ only in the numerically coded values of those genes. You could go from any creature in the safari park to any other, or to any of trillions of other biomorphs, simply by selective breeding.

Figure 1.15 Safari Park of Biomorphs, bred with the 'Blind Watchmaker' computer program.

The most recent version of the program can breed biomorphs that vary also in colour. It is based on the old program but it has a more elaborate 'embryology' and new genes controlling the colour of tree branches. There are also new genes determining whether each branch of the tree is a line or a rectangle or an oval, whether the shape concerned is filled or empty, and how thickly the lines are painted. Figure 1.16 doesn't convey how stunning these coloured biomorphs can be because, for reasons of economy, it is printed not in colour but in grey and white. When using the colour program I find myself following evolutionary alleys not towards insects and scorpions but towards flowers and the sort of abstract patterns that might look good on wallpaper or bathroom tiles. Figure 1.16 Safari park of biomorphs bred by 'Colour Watchmaker', here displayed in grey scale for economic reasons. The large black and white triangles in the background were added purely for decorative reasons. They were not part of the original breeding environment.

Biomorphs are 'artificially selected' by a human chooser. In this respect they are like cabbages or pedigree dogs. But artificial selection requires a human chooser and is not the main subject of this book. Following Darwin himself, I am using artificial selection as a model for a different process: natural selection. Finally, the time has come to introduce natural selection itself. Natural selection is like artificial selection, but without the human chooser. Instead of a human deciding which offspring shall die and which shall reproduce, nature 'decides.' The quotation marks are vital because nature doesn't consciously decide. This might seem too obvious to emphasise, but you'd be surprised by the number of people who think natural selection implies some kind of personal choice. They couldn't be more wrong. It just is the case that some offspring are

more likely to die while others have what it takes to survive and reproduce. Therefore, as the generations go by, the average, typical creature in the population becomes ever better at the arts of surviving and reproducing. Ever better, I should specify, when measured against some absolute standard. Not necessarily more effective in practice because survival is continually menaced by other creatures who are also evolving and perfecting their arts. A species may get progressively better at the art of avoiding predators but, since predators are concurrently getting better at the art of catching prey, there may be no net gain. This kind of 'evolutionary arms race' is fascinating, but we are jumping ahead of ourselves.

Artificial selection is relatively easy to achieve in the computer, and the biomorphs are a good example. It is my dream to simulate natural selection in the computer too. Ideally I'd like to set up the conditions for evolutionary arms races in which 'predators' and 'prey' would emerge on the screen and goad each other into progressive evolution while we sat back and watched. Unfortunately it is very difficult, for the following reason. I said that some offspring are more likely to die, and it might seem easy enough to simulate nonrandom death. But, in order to be a good simulation of a natural death, the demise of the computer creature must result from some interesting imperfection, like having short legs which make it run more slowly than predators. Computer biomorphs, for instance the insect-like forms in Figure 1.15, sometimes have appendages which we imagine we see as legs. But they don't use these 'legs' for anything, and they don't have predators. They don't have prey or food plants. There is no weather in their world and no disease. In theory we can simulate any of these hazards. But to model any one of them in isolation would be scarcely less artificial than artificial selection itself. We'd have to do something like arbitrarily decide that long, thin biomorphs can run away from predators better than short fat ones. It is not difficult to tell the computer to measure the dimensions of biomorphs and choose the lankiest for breeding. But the resulting evolution

would not be very interesting. We'd just see biomorphs becoming more and more spindly as the generations go by. It is no more than we could have achieved by artificially selecting the spindliest by eye. It does not have the emergent qualities of natural selection, which a good simulation might achieve.

Real life natural selection is much subtler. It is also in one sense much more complicated though in another sense it is deeply simple. One thing to say is that improvement along any one dimension, like leg length, is only improvement within limits. In real life there is such a thing, for a leg, as being too long. Long legs are more vulnerable to breaking and to getting tangled up in the undergrowth. With a little ingenuity, we could program analogues of both breakages and entanglements into the computer. We could build in some fracture physics: find a way of representing stress lines, tensile strengths, coefficients of elasticity - anything can be simulated if you know how it works. The problem comes with all the things that we don't know about or haven't thought of, and that means almost everything. Not only is the optimal leg length influenced by innumerable effects that we haven't thought of. Worse, length is only one of countless aspects of an animal's legs that interact with each other, and with lots of other things, to influence its survival. There is leg thickness, rigidity, brittleness, weight to carry around, number of leg joints, number of legs, taperingness of legs. And we've only considered legs. All the other bits of the animal interact with each other to influence the animal's probability of surviving.

As long as we try to add up all the contributions to an animal's survival theoretically, in a computer, the programmer is going to have to make arbitrary, human decisions, which mean that our simulation is scarcely an improvement over artificial selection by human eye. What we ideally should do is simulate a complete physics and a complete ecology, with simulated predators, simulated prey, simulated plants and simulated parasites. All these model

creatures must themselves be capable of evolving. The easiest way to avoid having to make artificial decisions might be to burst out of the computer altogether and build our artificial creatures as three-dimensional robots, chasing each other around a three-dimensional real world. But then it might end up cheaper to scrap the computer altogether and look at real animals in the real world, thereby coming back to our starting point! This point is less frivolous than it seems. I'll return to it in a later chapter. Meanwhile, there is a little more we can do in a computer, but not with biomorphs.

One of the main things that makes biomorphs so unamenable to natural selection is that they are made of fluorescent pixels on a two-dimensional screen. This 2-dimensional world doesn't lend itself to the physics of real life in most respects. Quantities like sharpness of teeth in predators and strength of armour plating in prey; quantities like muscular strength to throw off a predator's attack or virulence of a poison do not emerge naturally in a world of 2-dimensional pixels. Can we think of a real-life case of, say, predators and prey, which does lend itself, naturally and without contrived artificiality, to simulation on a 2-dimensional screen? Fortunately we can. I've already mentioned spider webs when talking about designoid traps. Spiders have 3-dimensional bodies and they live in a complex world of normal physics like most animals. But there is one particular thing about the way some spiders hunt that is peculiarly suited to simulating in two dimensions. A typical orb web is, to all intents and purposes, a 2-dimensional structure. The insects that it catches move in the third dimension, but at the moment of truth, when an insect is caught or escapes, the action is all in one 2-dimension plane, the plane of the web. The spider web is as good a candidate as I can think of for an interesting simulation of natural selection on a 2-dimensional computer screen. The next chapter is largely devoted to the fascinating story of spider webs, beginning with the natural history of real webs and moving on to computer models of webs and their evolution by 'natural' selection in the computer.

